

Bayesian analysis approaches to risk modeling

Bob Youngblood, Idaho National Laboratory

Robert.Youngblood@inl.gov

Meeting of PHMSA Risk Model Work Group


August 9-11, 2016

www.inl.gov



Frequentists Vs. Bayesians

Xkcd.com/1132/



Frequentists Vs. Bayesians

Xkcd.com/1132/

DID THE SUN JUST EXPLODE?
(IT'S NIGHT, SO WE'RE NOT SURE.)

THIS NEUTRINO DETECTOR MEASURES
WHETHER THE SUN HAS GONE NOVA.

THEN, IT ROLLS TWO DICE. IF THEY
BOTH COME UP SIX, IT LIES TO US.
OTHERWISE, IT TELLS THE TRUTH.

LET'S TRY.

DETECTOR! HAS THE
SUN GONE NOVA?

(ROLL)
YES.

FREQUENTIST STATISTICIAN:

THE PROBABILITY OF THIS RESULT
HAPPENING BY CHANCE IS $\frac{1}{36} = 0.027$.

SINCE $p < 0.05$, I CONCLUDE
THAT THE SUN HAS EXPLODED.

Frequentists Vs. Bayesians

Xkcd.com/1132/

Remember to
use all of your
information.

DID THE SUN JUST EXPLODE?
(IT'S NIGHT, SO WE'RE NOT SURE.)

THIS NEUTRINO DETECTOR MEASURES
WHETHER THE SUN HAS GONE NOVA.

THEN, IT ROLLS TWO DICE. IF THEY
BOTH COME UP SIX, IT LIES TO US.
OTHERWISE, IT TELLS THE TRUTH.

LET'S TRY.

DETECTOR! HAS THE
SUN GONE NOVA?

ROLL
YES.

FREQUENTIST STATISTICIAN:

THE PROBABILITY OF THIS RESULT
HAPPENING BY CHANCE IS $\frac{1}{36} = 0.027$.
SINCE $p < 0.05$, I CONCLUDE
THAT THE SUN HAS EXPLODED.

BAYESIAN STATISTICIAN:

BET YOU \$50
IT HASN'T.

From a PHMSA notice of proposed rulemaking:

PHMSA proposes to clarify the risk assessment aspects of the IM rule to explicitly articulate functional requirements and to assure that risk assessments are adequate to: (1) evaluate the effects of interacting threats, (2) determine intervals for continual integrity reassessments, (3) determine additional preventive and mitigative measures needed, (4) analyze how a potential failure could affect HCAs, including the consequences of the entire worst-case incident scenario from initial failure to incident termination, (5) identify the contribution to risk of each risk factor, or each unique combination of risk factors that interact or simultaneously contribute to risk at a common location, (6) account and compensate for uncertainties in the model and the data used in the risk assessment, and (7) evaluate predicted risk reduction associated with preventive and mitigative measures.

From a notice of proposed rulemaking (continued):

While PHMSA does not propose to prescribe the specific risk assessment model operators must use, PHMSA does propose to clarify the characteristics of a mature risk assessment program. These include: (1) identifying risk drivers; (2) evaluating interactive threats; (3) assuring the use of traceable and verifiable information and data; **(4) accounting for uncertainties in the risk model and the data used; (5) incorporating a root cause analysis of past incidents; (6) validating the risk model in light of incident, leak and failure history and other historical information;** (7) using the risk assessment to establish criteria for acceptable risk levels; and (8) determining what additional preventive and mitigative measures are needed to achieve risk reduction goals. PHMSA proposes to clarify that the risk assessment method selected by the operator must be capable of successfully performing these functions.

Purpose

- Of the RMWG:
 - To characterize the state of the art of pipeline risk modeling, ..., Identify a range of state of the art methods & tools ..., provide recommendations to PHMSA regarding the use of these methods, tools, and data requirements.
- Of this talk:
 - To help the RMWG do that, in the area of Bayesian analysis of operating experience data, in light of the PHMSA needs quoted above

Outline

- Selected highlights of “Decisions Under Uncertainty”
 - World War II: Scramble, or not?
 - False positives / false negatives
 - Value of Information

Theme: Dealing with uncertainty in decision-making
- Bayesian Analysis
 - What does it do?
 - Ways to go wrong with it: Hypothesis space, priors, “likelihood” models
 - ESP example
 - Data example

Theme: How we analyze uncertainty
- Selected highlights of how NRC and industry handle data on nuclear power plants
 - What sorts of decisions are supported, and how

A Fundamental Practical Problem

- We need to decide whether to mitigate an adverse condition (hazard, disease, performance issue)
- Different versions of the problem: either
 - we are not sure whether the condition exists, and we are dealing with Prob (it exists); or
 - Adverse consequences occur with probability p , but we are not sure about the value of p
- We have some information, but it does not completely determine the “state of nature”
- Being wrong in the assessment has significant adverse consequences
- Questions:
 - What is the best decision we can make, given the information we have?
 - How much would more definitive information be worth?
 - Should we expend the time and resources to get that information before choosing whether to mitigate?

Next slides

- Discuss Two-alternative model (To mitigate, or not to mitigate)
 - Each alternative has different possible outcomes
 - Probabilities and Consequences are given
 - Show the “Decision Tree”
 - Discuss pros and cons of the two possible decisions
- Illustrate “value of information:”
 - what it could be worth to obtain evidence that will make the decision more obvious by reducing uncertainties
 - Expand the decision tree to reflect not only the choice among the original alternatives, but also the decision whether to gather more information

Simple Version of the Problem

- A hazard (potential for adverse consequences) has been identified
- Steps could be taken to deal with the hazard
- Need to decide whether to take the steps (implement the system, ...)
- Given:
 - Probability of scenario leading to adverse consequences
 - Consequences (in monetary units) conditional on that scenario
 - Cost (in monetary units) and effectiveness of the mitigating system
- Pros and Cons:
 - Taking steps largely mitigates the hazard, but costs money
 - Not taking steps saves the cost of mitigation, but may incur the adverse consequences

Stakes Associated with Incorrect Diagnoses

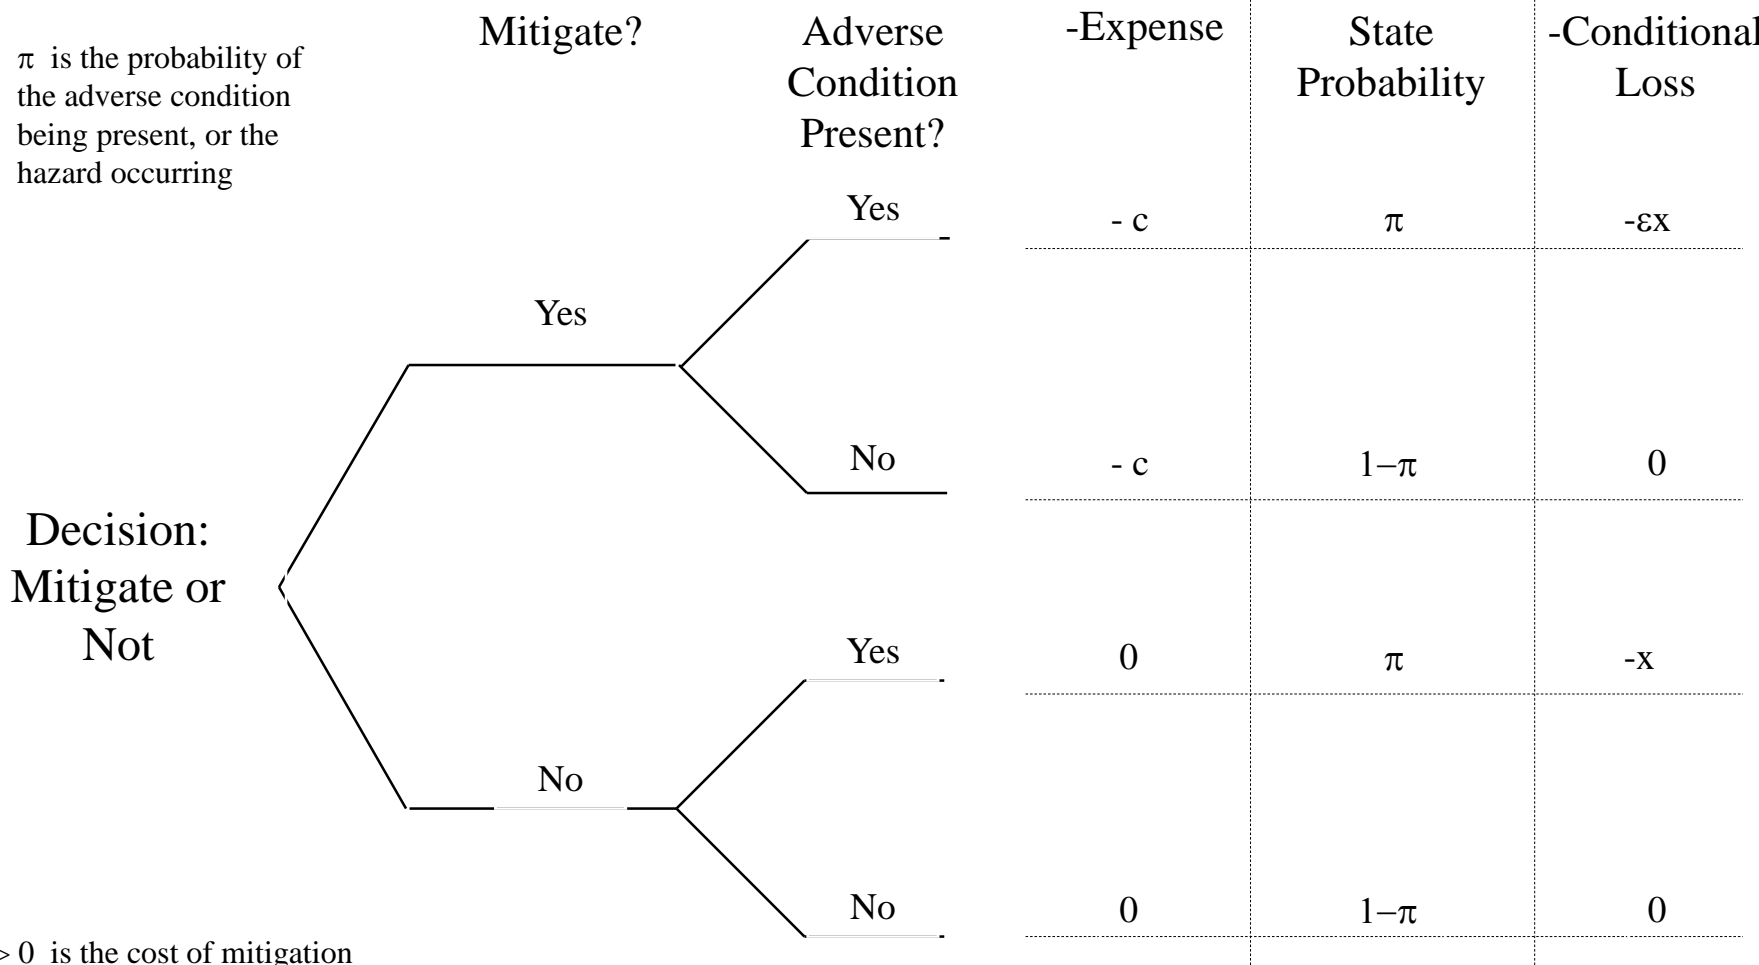
Terminology:

- False positive: conclusion that adverse condition is present, when it is not
- False negative: conclusion that adverse condition is not present, when it is

| <u>Context</u> | <u>Consequences of False Indications</u> |
|---|---|
| Medical Diagnosis (or Non-Destructive Evaluation) | <p>False positive: treat for disease (or replace component), unnecessarily incurring costs and side effects (or rejecting a “good” item)</p> <p>False negative: leave disease untreated (leave flawed component in place)</p> |
| Adding a mitigating system | <p>False positive: Incur cost</p> <p>False negative: Experience unmitigated consequences</p> |
| Enemy Aircraft Detection | <p>False positive: scramble for no reason</p> <p>False negative: fail to defend against attack</p> |

States & Outcomes of Mitigation Decision

π is the probability of the adverse condition being present, or the hazard occurring



$c > 0$ is the cost of mitigation

$\epsilon > 0$ but $\ll 1$ is introduced to allow for the possibility that mitigation is not perfect.

x is the (monetized) cost resulting from the hazardous event without mitigation

Decision based on expected consequences

- Expected consequences of “spend” (take the steps, mitigate, ...) decision:
 - Certain expenditure of “c”, experience consequences εx with probability π
 - $-\pi^*(-c-\varepsilon x)+(1-\pi)^*(-c) = -c-\pi\varepsilon x$
- Expected consequences of “don’t spend” decision:
 - Avoid cost “c” of mitigating system, but experience consequences x with probability π
 - $\pi^*(-x)$
- Decide to “spend” when
 - $[(-c-\pi\varepsilon x)]-[\pi^*(-x)] > 0$, or $\pi x(1-\varepsilon)>c$
 - That is, spend to mitigate when expected reduction in damages, $(\pi x(1-\varepsilon))$, exceeds cost c of mitigation
 - *Note: this simple formalism contemplates monetization of consequences, but does not address broader issues of utility such as risk aversion.*

What would it be worth to eliminate uncertainty?

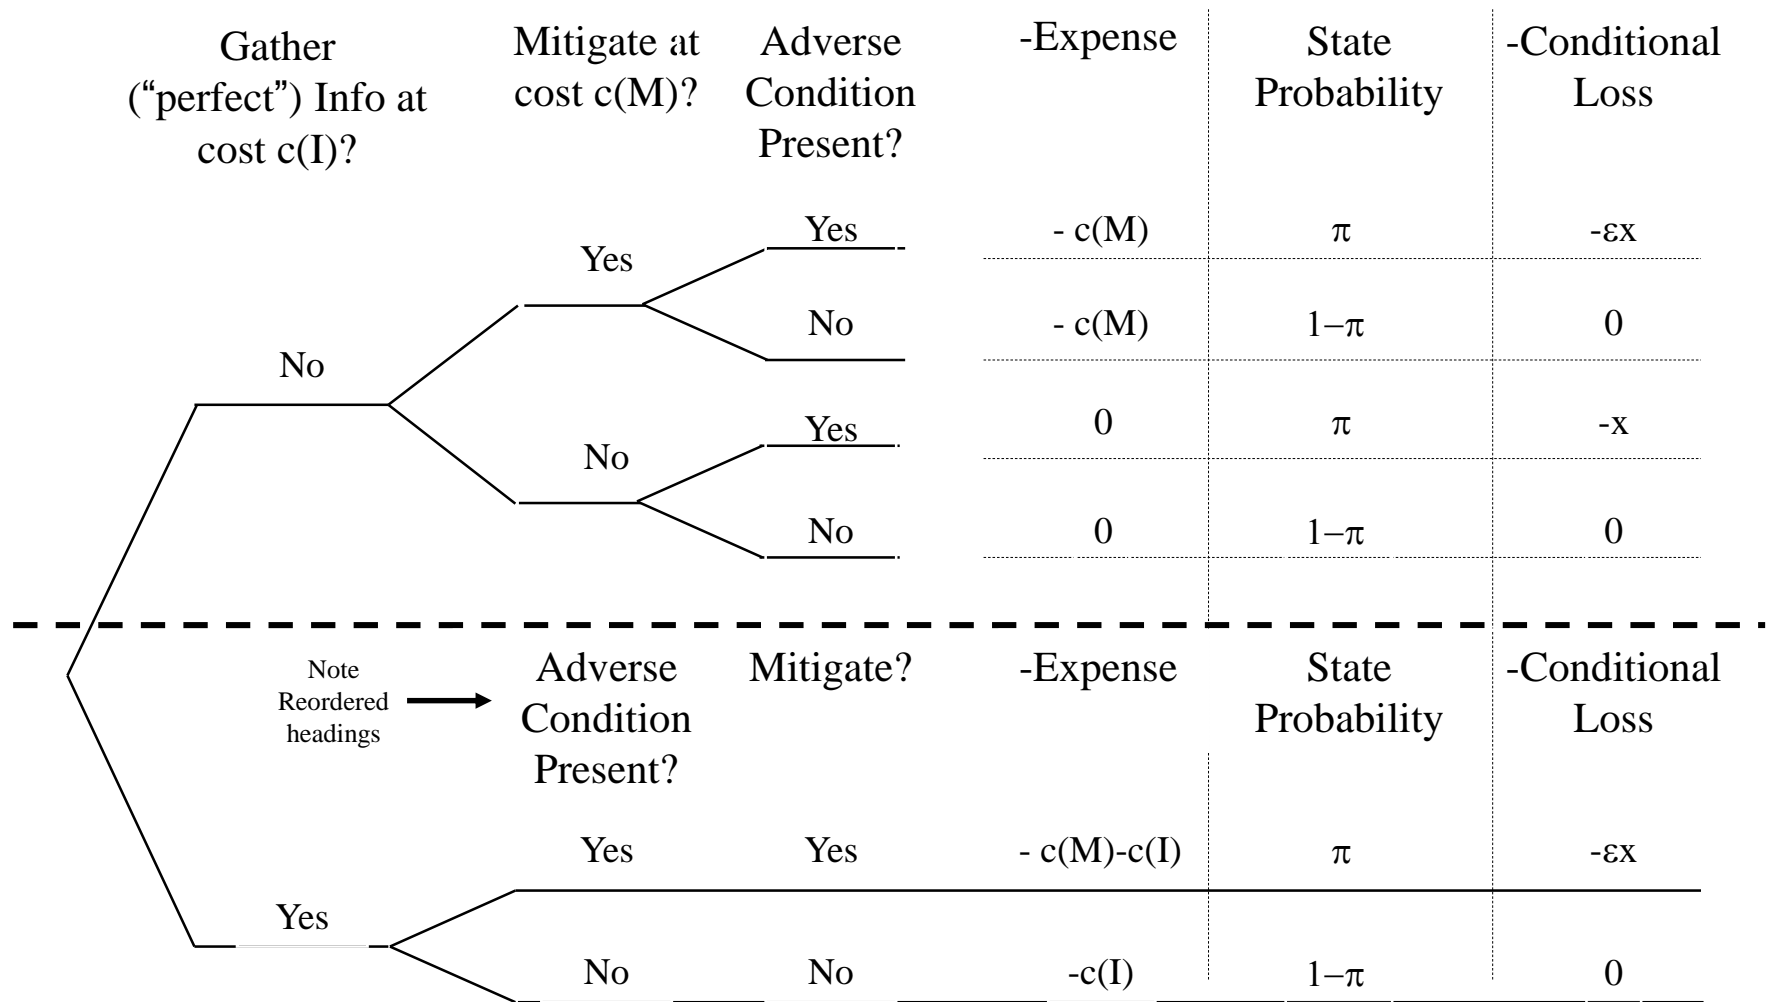
- In the example, the hazard is not known to exist: it has probability π
- We will improve the expected consequences of our decision if we KNOW whether the hazard is present
- Suppose c , ε , x are such that mitigation is well worth while if the hazard is definitely present. Then:
 - If we know the hazard is present, we will spend to mitigate it
 - If we know the hazard is not present, we will save the money
- The Value of Perfect Information is the difference between
 - the expected consequences conditional on knowing, and
 - the expected consequences conditional on the current state of knowledge (uncertainty)

$$\text{VOPI} = \langle U \rangle_{\text{no uncertainty}} - \langle U \rangle_{\text{uncertainty}}$$

Value of Information (VOI) for this example

- Expected consequences conditional on knowing:
 - $\pi \{ \text{consequences of best decision conditional on hazard being present} \} +$
 - $(1 - \pi) \{ \text{consequences of best decision conditional on hazard NOT being present} \}$
- This is $\pi * (-c - \epsilon X) + (1 - \pi) * 0$, or $\pi * (-c - \epsilon X)$
- Suppose the best decision under uncertainty is to “spend” (but there is some probability that we are wasting money)
 - Then the VOI is
 - $\square \pi * (-c - \epsilon X) - [-c - \pi \epsilon X] = c(1 - p)$
 - Amount wasted if no need, times the probability of no need
- Suppose the best decision under uncertainty is not to spend (but we may suffer unmitigated adverse consequences)
 - Then the VOI is
 - $\square \pi (-c - \epsilon X) - [\pi(-X)]$ or $\pi[X(1 - e) - c]$
 - Difference in damages that it would have made to spend, if hazard is present, times the probability that the hazard is present
- If the cost of mitigation exceeds the conditional consequences, the best decision is not to spend. Then the information has no value in this decision – VOPI = 0 - because information cannot change the decision (not to spend)


Gather [“perfect”] Information, Or Not?



Example: Best to Get Information

Fairly low
probability, high
consequence


Conditional
Consequences $x = \$1\text{E}7$
 $\pi = 0.1$
 $c(M) = \$1\text{E}6$
 $c(I) = \$1\text{E}5$



Example: Best just to mitigate

Fairly high
probability, high
cost of
information

Conditional
Consequences $x = \$1\text{E}7$
 $\pi=0.5$
 $c(M)=\$1\text{E}6$
 $c(I)=\$6\text{E}5$




Example: Mitigation Costs > Consequences

Cost of Mitigation is prohibitive (so don't mitigate);

Information cannot change the decision;
VOPI=0

Conditional
Consequences $x = \$9E5$
 $\pi=0.5$
 $c(M)=\$1E6$
 $c(I)=\$6E5$



Summary of Two-State Example

- Caveats:
 - This example is based on minimizing monetary losses
 - There is no consideration of risk tolerance, or utility other than monetary loss
- Uncertainty can limit the expected utility associated with a given decision situation
- Analysis can show how much it is potentially worth to reduce uncertainty
- In such a case, the uncertainty has a financial implication

The “value of perfect information” example just shows what it could be worth to reduce uncertainty. So: we’ve done a real-world measurement, and have some information (not perfect information).

Now what?

=>Bayesian analysis.

“Aleatory” vs. “Epistemic”

- Epistemic: State of knowledge uncertainty
- Aleatory: Variability from one trial to the next
- Example of Aleatory uncertainty:
 - We **know** Mean Time Between Failures (MTBF) for a component type, but we don't know when any specific component of that type will fail
- Example of Epistemic uncertainty:
 - We **DON'T know** MTBF
- In general, in this talk, we are talking about epistemic uncertainty: uncertainty in a model parameter
 - Could be reliability models (failure rates, failure probabilities, equipment availabilities)
 - Could be physical model parameters (toughness, corrosion rates, ...)

What does Bayesian analysis do?

- It shows us how to incorporate newly acquired evidence into our current state of knowledge regarding some parameter. Examples:
 - What does recent operating experience tell us about the failure rates of components in our system?
 - We thought the compressor failure rate was λ , but based on that, we should have had only n failures; and instead we've had $m > n$ failures.
 - What do recent test results tell us about the parameters of physical models, or even the applicability of those models to our situation?

Bayes' Theorem:

- Bayes' "theorem" states that

$$p(H_i | E) = P(H_i) \frac{p(E | H_i)}{p(E)},$$

$$p(E) = \sum_i p(E | H_i) p(H_i)$$

- where
 - H_i represents a hypothesis whose probability is to be updated with new evidence,
 - $p(H_i)$ is the prior probability of H_i ,
 - E represents a new piece of evidence,
 - $p(x|y)$ is the conditional probability of x given y ,
 - $p(E)$, the prior probability of the observed evidence

Bayes' Theorem:

- Bayes' "theorem" states that

What we think now

What we used to think

$$p(H_i | E) = P(H_i) \frac{p(E | H_i)}{p(E)},$$

Factor measuring the consistency of the observed evidence E with the various competing hypotheses H_i

$$p(E) = \sum_i p(E | H_i) p(H_i)$$

- where
 - H_i represents a hypothesis whose probability is to be updated with new evidence,
 - $p(H_i)$ is the prior probability of H_i ,
 - E represents a new piece of evidence,
 - $p(x|y)$ is the conditional probability of x given y,
 - $p(E)$, the prior probability of the observed evidence

Application of Bayes' Theorem:

Gorillas in the room

- Structure of the hypothesis space: what ARE the $\{H\}$?
 - As in the “ESP example” (later slides)
 - Other examples abound
- Selection (formulation) of the prior $p(H)$
 - Has a huge effect on the results
 - Still a research topic
- Modeling of the likelihood $p(E|H_i)$
 - If you're comparing a model to data, then this includes all sorts of things
 - Model form uncertainty
 - Selection of data (E): which data apply?

Really only one gorilla

- Bayes' "theorem" states that

$$p(H_i | E) = P(H_i) \frac{p(E | H_i)}{p(E)},$$

$$p(E) = \sum_i p(E | H_i) p(H_i)$$

- Everything on the right-hand side includes modeling choices made by the user
- So the "theorem" is an identity, but you can still go very wrong
- We have met the enemy, and he is us

Structure of the hypothesis space

(Probability Theory: The Logic of Science, E. T. Jaynes)

- Consider an experiment to determine whether an individual has extrasensory perception (ESP)
 - Experiment involves seeing whether individual can sense which of several possible cards is held
- It's possible to guess correctly sometimes, but (barring ESP) EXTREMELY unlikely to guess correctly a large fraction of the time
- Consider two hypotheses: yes (ESP) and no (no ESP)
- Establish prior probabilities for these two hypotheses
- The data come in. The individual gets everything right. You update your prior with data, and it looks like the individual has ESP.
- **What do you conclude?**

Brief Digression on odds ratio

- Interesting to focus on the ratios

- Write Bayes' formula for each hypothesis;

$$\left\{ \begin{array}{l} p(Yes | E) = P(Yes) \frac{p(E | Yes)}{p(E)}, \\ p(No | E) = P(No) \frac{p(E | No)}{p(E)}, \end{array} \right.$$

- divide one by the other.

$$\frac{p(Yes | E)}{p(No | E)} = \frac{p(Yes)}{p(No)} * \frac{p(E | Yes)}{p(E | No)}$$

Brief Digression on odds ratio

$$\frac{p(Yes | E)}{p(No | E)} = \frac{p(Yes)}{p(No)} * \frac{p(E | Yes)}{p(E | No)}$$

Strength of the Prior

Strength of the Evidence

- The factor on the right can be meaningfully discussed, independently of the prior
- Expert can say “probability of E (this DNA result) given “Yes” (e.g., that the accused is guilty) is 10000 times the probability of this DNA result given “No” (that the accused is innocent).”
- Just involves fingerprint technology and the chain of evidence, no knowledge of the accused or other evidence.

Structure of the hypothesis space

(Probability Theory: The Logic of Science, E. T. Jaynes)

- At this point, Jaynes would admit a third hypothesis: deception
- The answer YOU get may not be the answer Jaynes gets
 - But his posterior will assign higher probability to “deception” than to ESP

1983 NRC Guidance on the prior

- 5.5.2 BAYESIAN ESTIMATION
- The Bayesian approach is similar to the classical approach in that it yields "best" point estimates and interval estimates, the intervals representing ranges in which, we are confident, the parameter really lies. It differs in both practical and philosophical aspects, though. The practical distinction is in the incorporation of belief and information beyond that contained in the observed data; the philosophical distinction lies in assigning a distribution that describes the analyst's belief about the values of the parameter. This is the so-called prior distribution.
- The prior distribution may reflect a purely subjective notion of probability, as in the case of a Bayesian degree-of-belief distribution, or any physically caused random variability in the parameter, or some combination of both.

- If uncertain parameters have to be estimated, and there is no practical alternative to using expert elicitation, consider the Kaplan “expert evidence” approach,* as opposed to simply asking experts to estimate uncertain parameters directly. Note, however, that the Kaplan idea stops at evidence-gathering, leaving the task of formulating a likelihood function yet to be done.
- Under some conditions, it may be appropriate to apply the Cooke method.
- *Note: The Kaplan method was formulated to avoid the earlier practice of treating expert opinions as if they were experimental results, which can easily lead to controversy. The Cooke method operates within the tradition of treating expert opinions analogously to experimental results, but does so in an intelligent way within a process that evaluates the experts with calibration questions that provide reasonably objective performance weights for experts. The Cooke method was not the first to weight experts, but it does so within an objective, algorithmic process.*

** Instead of asking them what they think the answer is, you ask them what evidence informs their assessment. You pool their evidence, not their answers; and then you process that pooled body of evidence using Bayes’ theorem.*

Example from NASA PRA Procedures Guide

Example 5: Updating Distribution of Failure on Demand Probability

It is assumed that the prior distribution of a component failure probability on demand is characterized by a beta distribution with Mean = $1\text{E-}4$ failures per demand, and Standard Deviation = $7\text{E-}5$. It is also assumed that the operational data for the component category indicate 1 failure in 2,000 demands. Since the prior distribution is a Beta, and the likelihood function is Binomial, the posterior distribution is also a Beta distribution. Both the prior and posterior distributions are shown in Figure 5-5.





Figure 5-5. The Prior and Posterior Distributions of Example 5.

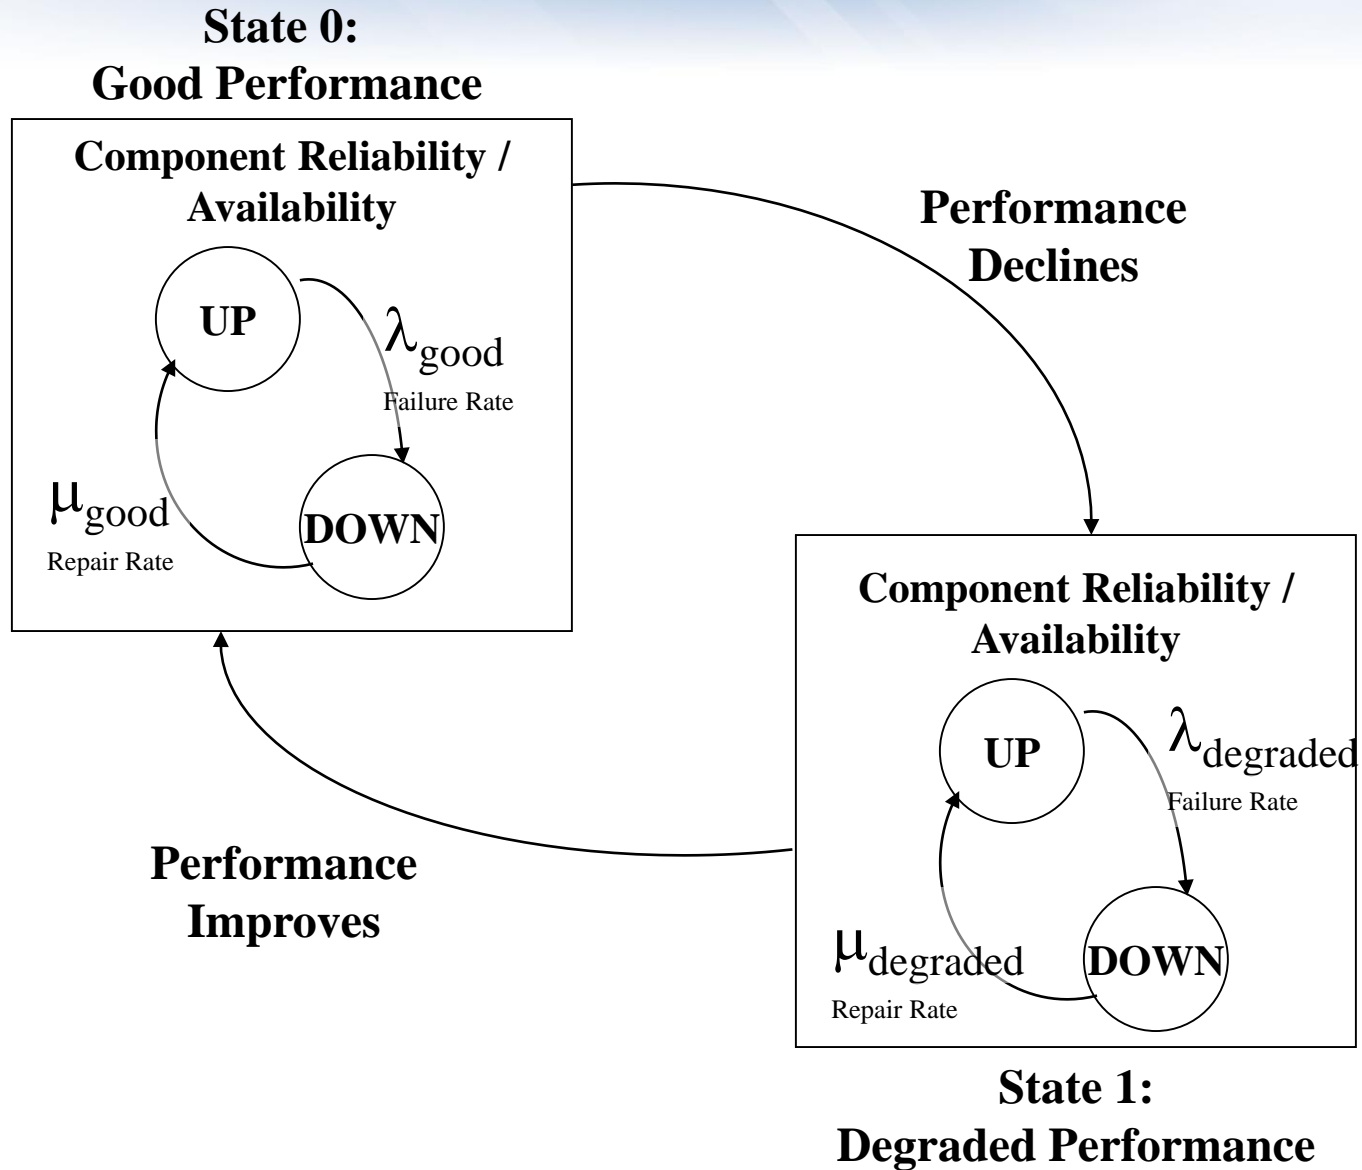
"Example 5" Revisited



Comments on Example


- The new data are not absolutely inconsistent with the prior (1/2000 could have been a fluke) but one should ask whether things have gotten worse than they were when the prior was formulated
- The Bayesian formalism is robust: it will give you an answer whether or not your underlying assumptions are right
 - Are all these trials “exchangeable?” Or has something changed (gotten worse) since the prior distribution was derived?
- Consider an application to performance assessment.
 - Specifically: Is performance getting worse? (Is unreliability increasing?)
 - Instead of a relatively narrow prior distribution, consider a sum of distributions: one for good performance, and one for bad performance.

Mixture Priors- 1




Mixture Priors- 2:


Compare the behavior of two prior distributions, *gmix* and the “constrained non-informative prior” (CNIP)



Mixture Priors- 4: Despite having the same prior mean, the posterior distributions behave very differently



Posterior $E(p)$ from Several Priors



Calculation of the “Mitigating Systems Performance Index” (MSPI) Unreliability Index (URI)

NUREG-1816
Appendix M

$$\Delta CDF \approx \left. \frac{\partial CDF}{\partial A} \right|_0 \Delta A + \frac{1}{2} \left. \frac{\partial^2 CDF}{\partial A^2} \right|_0 (\Delta A)^2$$

\approx

$$= B\Delta A + \frac{1}{2} B' \cdot (\Delta A)^2$$

$$URI = \sum_{j=1}^m \left[\begin{aligned} &B_{Dj}(UR_{DBCj} - UR_{DBLj}) \\ &+ B_{Lj}(UR_{LBCj} - UR_{LBLj}) \\ &+ B_{Rj}(UR_{RBCj} - UR_{RBLj}) \end{aligned} \right]$$


j indexes failure modes for a particular component (in this case, a diesel generator)

NEI 99-02 Rev 6

- B_{Dj} , B_{Lj} and B_{Rj} are the Birnbaum importance measures for the failure modes fail on demand, fail to load, and fail to run respectively,
- UR_{DBC} , UR_{LBC} , and UR_{RBC} are Bayesian corrected plant specific values of unreliability for the failure modes fail on demand, fail to load, and fail to run respectively, and
- UR_{DBL} , UR_{LBL} , and UR_{RBL} are Baseline values of unreliability for the failure modes fail on demand, fail to load, and fail to run respectively.


NRC Site Info on MSPI for a particular plant

Mitigating Systems Performance Index, Emergency AC Power System



Thresholds: White > 1.00E-6 Yellow > 1.00E-5 Red > 1.00E-4

Mixture Priors- Vary the prior probability of degraded performance



Different Priors, Different Decisions!

| No. failures | CNI | | Logistic-Cauchy | | Robust Bayes | | Imprecise Probability | |
|--------------|-------|--------|-----------------|-------|--------------|-------|-----------------------|----------------|
| | Green | White | Green | White | Green | White | Green | White |
| 5 | 90.8 | -181.6 | 22.5 | -45.0 | 46.8 | -93.6 | (44.3, 100) | (-200, -88.6) |
| 6 | 81.4 | -162.8 | -4.8 | 9.7 | 22.0 | -44.0 | (19.3, 100) | (-200, -38.7)) |
| 7 | 67.6 | -135.1 | -25.5 | 51.0 | -1.4 | 2.8 | (-3.5, 99.9) | (-200, 6.9) |
| 8 | 50.1 | -100.2 | -37.1 | 74.3 | -19.5 | 38.9 | (-21.5, 99.8) | (-200, 43.0) |
| 9 | 30.6 | -61.1 | -44.6 | 89.1 | -32.8 | 65.5 | (-34.0, 99.5) | (-199, 68.0) |
| 10 | 11.1 | -22.2 | -47.7 | 95.5 | -41.9 | 82.0 | (-41.7, 98.9) | (-198, 83.5) |
| 11 | -6.4 | 12.7 | -49.2 | 98.4 | -45.6 | 91.3 | (-46.1, 97.8) | (-196, 92.1) |
| ... | ... | ... | ... | ... | ... | ... | ... | ... |
| 21 | ... | ... | ... | ... | ... | ... | (-50, 9.2) | (-18.4, 100) |
| 22 | ... | ... | ... | ... | ... | ... | (-50, -3.1) | (6.1, 100) |

Which prior to use??


At this point, we've illustrated each of the bullet items below

Application of Bayes' Theorem: Gorillas in the room

- Structure of the hypothesis space: what ARE the $\{H\}$?
 - As in the “ESP example” (later slides)
 - Other examples abound
- Selection (formulation) of the prior $p(H)$
 - Has a huge effect on the results
 - Still a research topic
- Modeling of the likelihood $p(E|H_i)$
 - If you're comparing a model to data, then this includes all sorts of things
 - Model form uncertainty
 - Selection of data (E): which data apply?

Application of data of different types

- The first example was for reliability (failures and demands)
- But the mixture prior example was, in part, for *performance states*
- If you can infer something about the performance state from inspection, that can be factored into the update as well
- Example on the following slide done assuming a significant false indication probability for “inspection”
 - Probability that inspection will conclude “degraded performance” even if performance is “good,” and *vice versa*



Formalism works for all kinds of things...



- Examples so far have stressed applications to reliability (failure rate, failure probability) based on evidence from operating experience (or “inspection”)
- But the Bayesian formalism works for all kinds of things ...
 - Subject of course to the caveats previously mentioned
- ... Such as parameters in physics models ...
 - ...Even complicated ones
 - ...Even many-parameter ones
 - ...Even hard-to-execute models, if you use Markov Chain Monte Carlo and model emulators

Forward vs. Backward Uncertainty Quantification (UQ)


Given the input distributions, what's the uncertainty in the prediction?



Input values of uncertain parameters, initial conditions, boundary conditions, etc...




Output metrics



How do the output distributions compare to observational data?

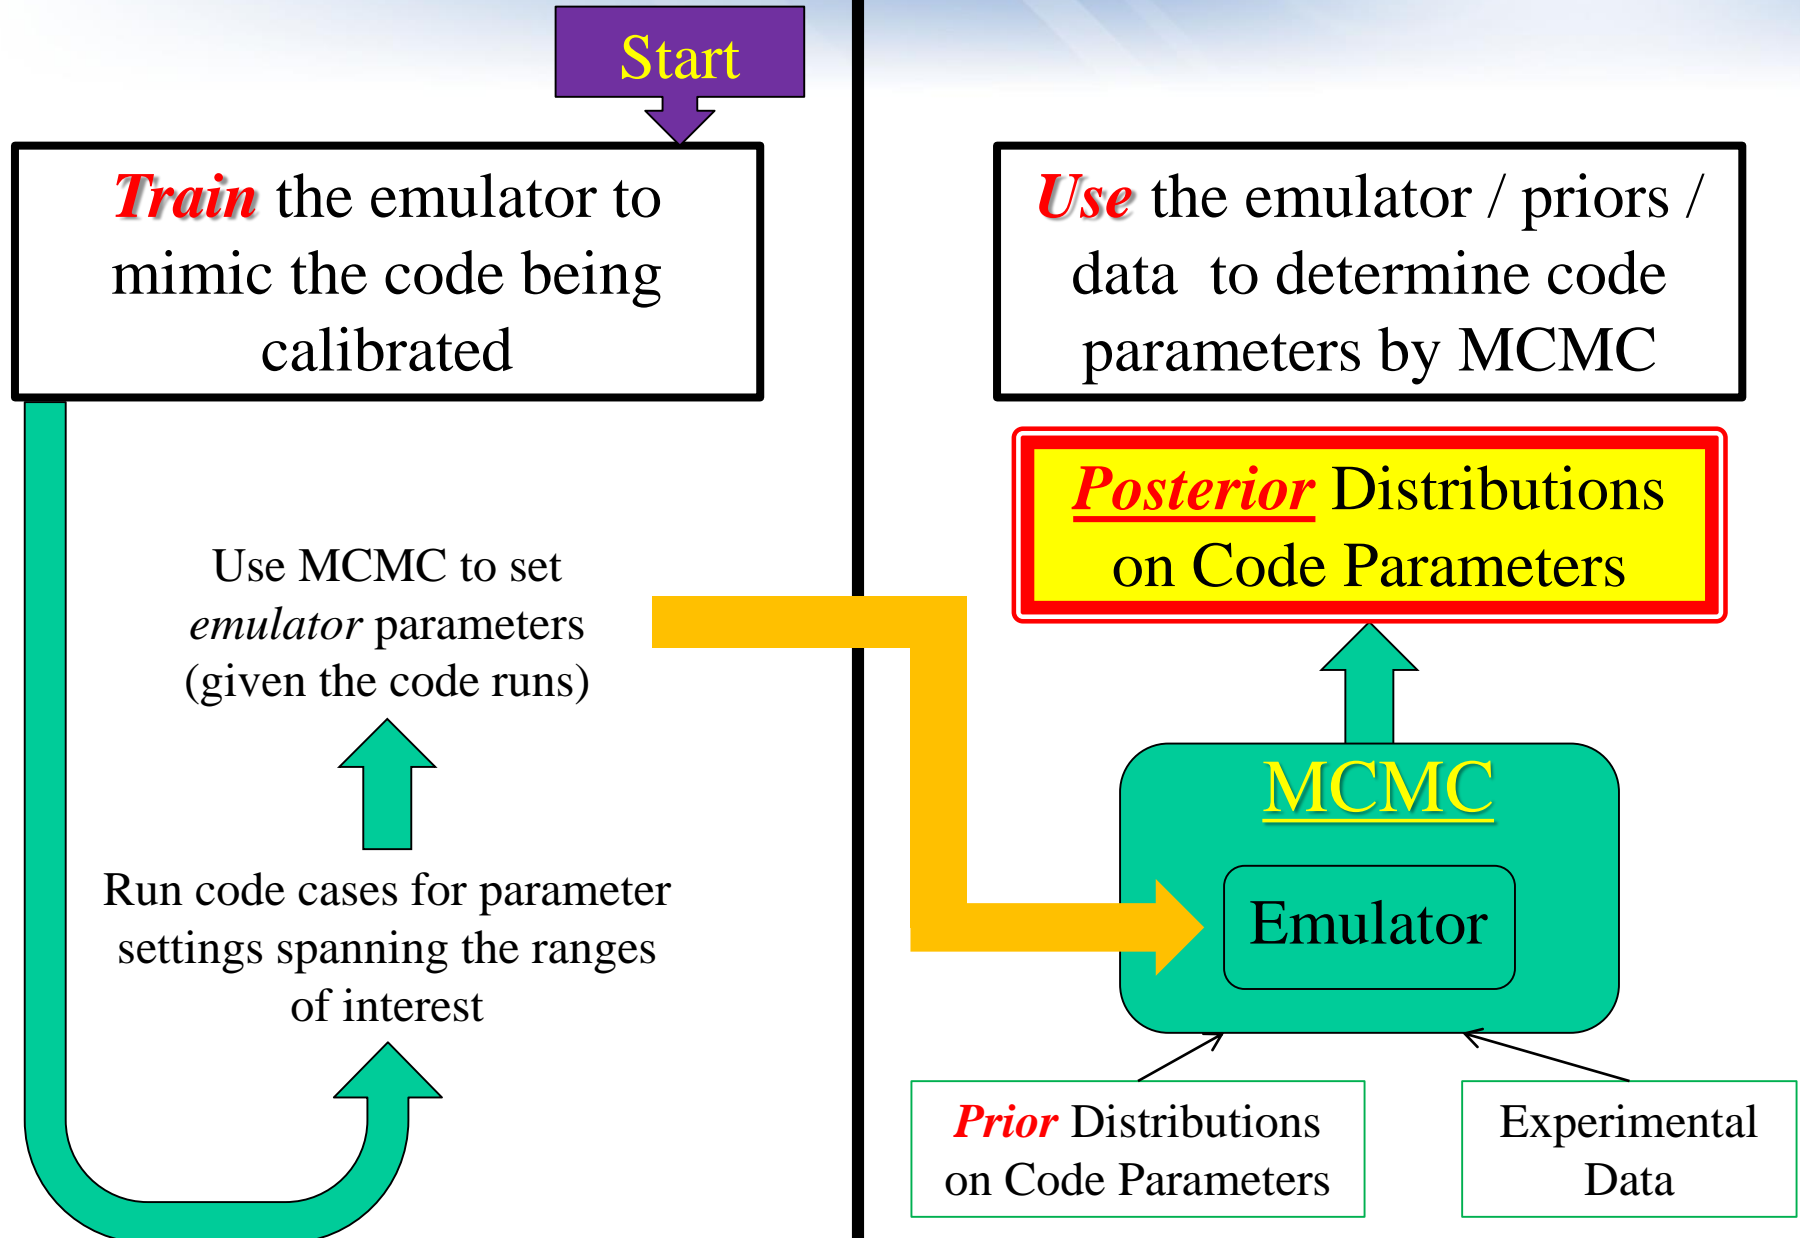


Experimental data



Given the experimental data, what's the joint distribution of the inputs?

Task: Estimate physical model parameters, given data



Complicated thermal-hydraulic model with lots of uncertain parameters, “calibrated” with experimental data using a Bayesian Markov Chain Monte Carlo approach.

The posterior predictions nail the observations.

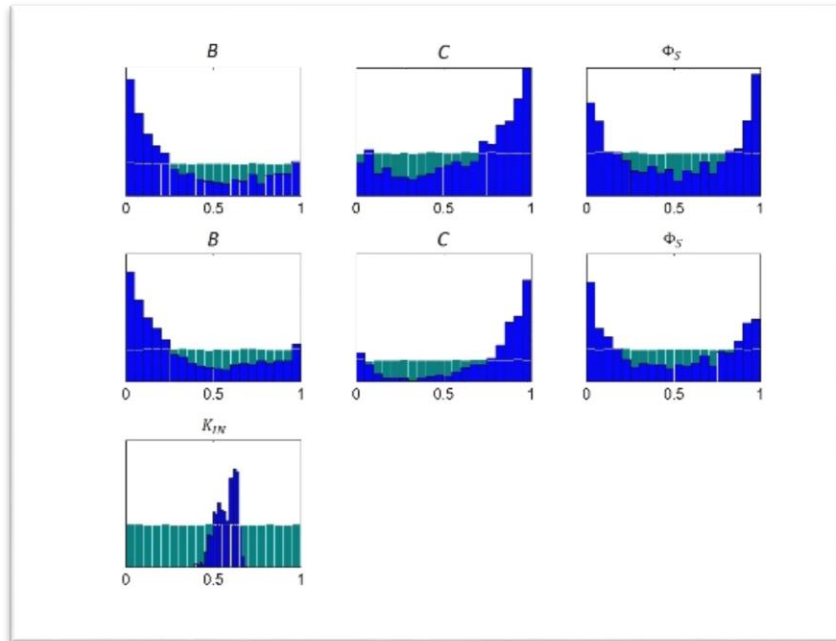


Fig. 9. IET only calibrated scaled posterior histograms




Fig. 8. IET calibrated posterior predictions relative to the “pseudo” data.

J.P. YURKO, *Uncertainty Quantification in Safety Codes Using a Bayesian Approach with Data from Separate and Integral Effect Tests*. Dissertation, MIT. Cambridge, MA, 2014.

Population Variability

Original idea: Kaplan, S. On a 'two-stage' Bayesian procedure for determining failure rates. *IEEE Transactions on Power Apparatus and Systems*, 1983, **PAS-102**, 195–262.

Bayesian parameter estimation in probabilistic risk assessment

Nathan O. Siu & Dana L. Kelly

Reliability Engineering and System Safety **62** (1998) 89–116
Published by Elsevier Science Limited
All rights reserved. Printed in Northern Ireland
0951–8320/98/\$19.00




Fig. 4. Illustrative plot of plant-specific failure probabilities.




Fig. 5. Hypothetical population variability curve.

The general idea:
Instead of pooling performance data from different sources (e.g., facilities), as if everybody's performance is the same: Develop a distribution expressing the variability in performance...




Fig. 5. Hypothetical population variability curve.

$$p(H_i | E) = P(H_i) \frac{p(E | H_i)}{p(E)},$$

$$p(E) = \sum_i p(E | H_i) p(H_i)$$

The general idea (continued):
... and use that distribution as a *prior* for the facility of current interest...

And update that prior with the **data** you have for the facility of current interest (“**E**”) to get a posterior distribution for the facility of current interest

This approach makes essential use of the idea that it makes sense to think in terms of family characteristics: that other facilities’ data carry implicit information about your facility.

General Principles:

- Strive to avoid the trap of understating uncertainty.
- Strive to make use of all available information that is legitimately applicable to the decision at hand.
- Maintain an essentially fallibilist posture with respect to analysis results.
- Be very careful about using the full standard Bayesian approach based on formulation and updating of an explicit prior.
 - If there is a lot of objective evidence to bring to bear, apply that evidence to a maximally ignorant prior, checking along the way to see whether the prior and the evidence are tugging the posterior in opposite directions.
 - “A lot of objective evidence” means “sufficient evidence that the posterior is reasonably insensitive to choice of prior.”
 - If data and prior are incompatible, ...

Summary

- It's extremely important to understand the uncertainties and what they do to the decision problem
- Bayes' theorem is a powerful tool for understanding the uncertainty, and for helping to figure out what to do in order to reduce it most effectively
 - Many problems in this arena might usefully map onto a “value of information” framework: what would it be worth to inspect / test / this pipeline?
 - *That question can be answered within classical decision analysis, if you understand your uncertainty.*
- A lot of theoretical capability has been developed.
- That capability has to be used with caution, because ...

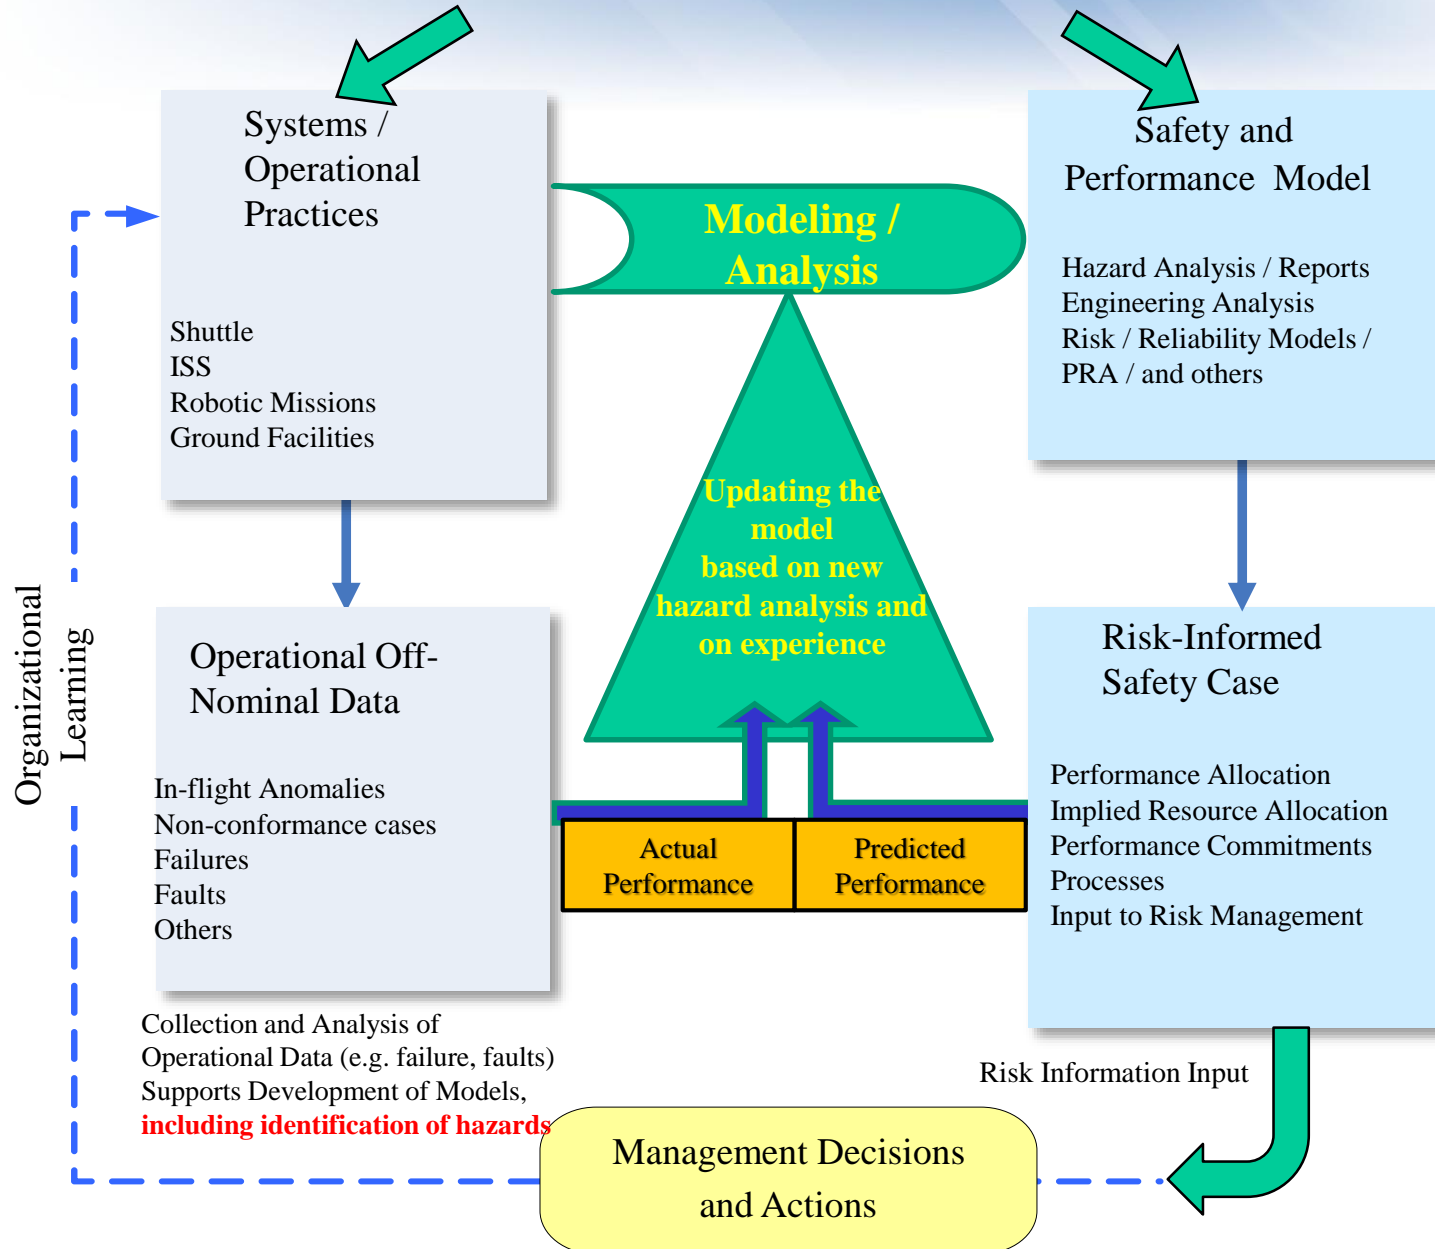
$$p(H_i | E) = P(H_i) \frac{p(E | H_i)}{p(E)},$$

$$p(E) = \sum_i p(E | H_i) p(H_i)$$

...this stuff is
all *user input*

- “Psychological Science Can Improve Diagnostic Decisions,” John A. Swets, Robyn M. Dawes, and John Monahan, *Psychological Science in the Public Interest* vol. 1, No. 1 (American Psychological Society, 2000).
- Robert T. Clemen, **Making Hard Decisions: An Introduction to Decision Analysis**, Second edition (Duxbury Press, 1996)
- **Probability Theory: The Logic of Science**, E. T. Jaynes (Cambridge University Press, 2003).
- **Information Theory, Inference, and Learning Algorithms**, David J.C. MacKay (Cambridge University Press, 2003).
- 'Expert information' versus 'expert opinions.' Another approach to the problem of eliciting/ combining/using expert knowledge in PRA,” S. Kaplan, *Reliability Engineering and System Safety* **35** (1992) pp. 61-72 (Elsevier, 1992). See also “Combining Probability Distributions from Experts in Risk Analysis,” S. Kaplan, *Risk Analysis* **20**, no. 2, pp. 155-156 (Society for risk Analysis, 2000).
- Summary from the epistemic uncertainty workshop: consensus amid diversity, Scott Ferson, Cliff A. Joslyn, Jon C. Helton, William L. Oberkampf, Kari Sentz, *Reliability Engineering and System Safety* **85** pp 355-369 (Elsevier, 2004).
- Bayesian parameter estimation in probabilistic risk assessment, Nathan O. Siu & Dana L. Kelly, *Reliability Engineering and System Safety* **62** (1998) 89-116.
- Effect of Epistemic Uncertainty Modeling Approach on Decision-Making: Example Using Equipment Performance Indicator, Dana Kelly and Bob Youngblood, *Proceedings of PSAM-11* (Helsinki, Finland, 2011).
- R. W. Youngblood and C. L. Atwood, “Mixture Priors for Bayesian performance monitoring 1: fixed-constituent model,” **Reliability Engineering and System Safety**, Volume 89, Issue 2, August 2005, Pages 151-163 (Elsevier, 2005).
- C. L. Atwood and R. W. Youngblood, “Mixture Priors for Bayesian performance monitoring 2: variable-constituent model,” **Reliability Engineering and System Safety**, Volume 89, Issue 2, August 2005, Pages 164-176 (Elsevier, 2005).
- C. L. Atwood and R. W. Youngblood, “Application of Mixture Priors to Assessment of Performance,” *Probabilistic Safety Assessment and Management (PSAM 7-ESREL '04)*, p. 444-449, edited by Cornelia Spitzer, Ulrich Schmocker, Vinh Dang (Springer, 2004).

Real World versus Models



Original version of this slide originated within NASA's Office of Safety and Mission Assurance. This version was probably oriented to accident precursor analysis.